

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 913 828 A2

(12)

EUROPEAN PATENT APPLICATION

- (43) Date of publication:
06.05.1999 Bulletin 1999/18
- (21) Application number: 98308387.4
- (22) Date of filing: 14.10.1998

(51) Int Cl. 6: G11C 7/00

- (84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI
- (30) Priority: 28.10.1997 US 959636
- (71) Applicant: MMC Networks, Inc.
Sunnyvale, CA 94086 (US)

- (72) Inventors:
 - Joffe, Alexander
Palo Alto, CA 94306 (US)
 - Birger, Ari
Sunnyvale, CA 94087 (US)

- (74) Representative: Gill, David Alan
W.P. Thompson & Co.,
Celcon House,
289-293 High Holborn
London WC1V 7HU (GB)

(54) Memory system and method of accessing the same

(57) In a memory system, each data bus (data) is connected to memories (110) connected to different address buses (Addr). Each memory (110) allows pipelined read operations such that when data are being read out from a memory (110) in one read operation, the address can be provided to the memory (110) for another read. However, write operations are not pipelined, and the write address and write data are provided to the memory simultaneously. Nevertheless, consecutive reads can overlap with writes. Each write operation uses address (Addr) and data (Data) buses not taken by any read occurring in parallel with the write. The address

(Addr) and data (Data) buses are connected to the memories so that no data bus penalty occurs when a memory is switched from a read to a write or from a write to a read. In some embodiments, multiple memories are subdivided into sets of mirror-image memories. In each set, all the memories store the same data. When simultaneous read accesses are desired to read data stored in one of the memories, the read accesses can be performed instead to different memories that are mirror images of each other. When any memory is written, all the memories of the same set are written with the same data.




FIG. 1

Description

[0001] The present invention relates to memories, systems using memories, and methods for accessing memories.

[0002] To increase throughput of memory access operations, some data processing systems employ multi-ported memories. Multiple access operations are allowed to proceed in parallel through different ports to increase the throughput. However, the cost of memories increases with the number of ports.

Therefore, it is desirable to use memories having fewer ports while still obtaining high throughput. Further, in memory systems with multi-ported memories, separate address and data buses are used for each port. It is desirable to reduce the number of address and data buses in the memory system.

[0003] It is also desirable to increase the address and data bus utilization in memories that use different timing for read and write operations. Such memories include fast synchronous SRAMs (static random access memories). Different timing for read and write operations causes address or data bus utilization penalty when the memory is switched from a write operation to a read operation or from a read to a write. It is desirable to reduce or eliminate such penalty.

[0004] It is also desirable to provide memory systems that enable one to obtain a non-blocking ATM (asynchronous transfer mode) switch, or some other switch, by combining two or more switch fabrics to increase the number of ports but without increasing the cost per port.

[0005] Some embodiments of the present invention provide memory systems that allow multiple access operations to proceed in parallel. Some embodiments use memories with different timing for read and write operations. In particular, in some embodiments, each memory in the memory system allows pipelined read access such that when data are being read out of the memory in one read operation, a read address is provided to the memory for another read operation. However, in a write operation, the address and data are provided to the memory at the same time. (Some fast synchronous SRAMs have such timing.) Therefore, when the memory is switched from a write to a read or from a read to a write, there is utilization penalty with respect to memory address or data ports. However, in some embodiments of the invention, no penalty occurs with respect to the data buses of the memory system. In some embodiments, no penalty occurs also with respect to the address buses. Further, the number of ports in each individual memory is reduced. In some embodiments, each memory is single-ported. The number of data buses is also reduced by making the data buses shared between different ports. In some embodiments, the number of address buses is also reduced by making them shared.

[0006] These advantages are achieved in some embodiments by connecting an address bus to different address ports corresponding to data ports connected to different data buses, and/or connecting a data bus to different data ports corresponding to address ports connected to different address buses. In some embodiments, each combination of an address bus and a data bus can be used to access a separate memory.

[0007] For example, Fig. 1 shows four single-ported memories 110_UL, 110_UR, 110_DL, 110_DR (Fig. 1). Address bus mAddr_U is connected to the address ports of memories 110_UR and 110_UL, and address bus mAddr_D is connected to the address ports of memories 110_DR and 110_DL. Data bus Data_L is connected to the data ports of memories 110_UL and 110_DL, and data bus Data_R is connected to the data ports of memories 110_UR and 110_DR.

[0008] Each combination of an address bus and a data bus allows access to one of the four memories 110.

[0009] The memory system is a shared memory in any flow switch, for example, an ATM (Asynchronous Transfer Mode) switch composed of two switch fabrics. The two address buses and the two data buses allow a write and a read to proceed in parallel. In each clock cycle, an ATM cell is written by one of the switch fabrics into one of the memories 110 for storage before transmission, and another cell is read by the other fabric from another one of the memories for transmission. A cell can be written into any memory available for a write operation. Each switch fabric has the same number of ports. Hence, when the switch fabrics are combined, the number of ports is doubled. However, the cost per port is about the same as in a single switch fabric.

[0010] Each memory is a synchronous SRAM allowing pipelined reads. A read operation latency is two clock cycles, with one cycle for the address and one cycle for the data. A write operation latency is one cycle. However, both of the address buses and both of the data buses can be used in every cycle. The bus utilization penalty is avoided as follows.

[0011] In each clock cycle, one address bus and one data bus can be taken by read operations. (More particularly, in each clock cycle, one address bus can carry an address for a read operation started in this cycle, and one data bus can carry data for a read operation started in the previous cycle). Therefore, in each clock cycle, one address bus and one data bus remain available for a write operation. No utilization penalty occurs.

[0012] The invention is not limited to two address buses or two data buses, or to four memories, or to any particular number of clock cycles needed for a read or a write.

[0013] Some embodiments provide mirror-image memories that duplicate each other. Some such embodiments allow multiple reads to occur in parallel. Thus, some embodiments are used in an ATM switch in which two reads and two writes can proceed in parallel. If two cells to be read out simultaneously are stored in the same memory, they are also stored in the mirror-image of that memory, and hence each cell can be read out from a different memory. Therefore, the two cells can be read out simultaneously even if each memory is single-ported.

[0014] When a cell is written into any memory, it is also written into the mirror image of that memory.

[0015] Some embodiments provide more than one mirror images for each memory. For example, in some embodiments, a switch can read four cells simultaneously, and each memory has three mirror image memories such that all the four memories store the same data. Each memory is single-ported. If more than one cells to be read out simultaneously are stored in the same memory, these cells are read out from different memories that are mirror images of each other.

5 Alternatively, in some embodiments each memory has only one mirror image, but each memory is double-ported, and hence four cells stored in the same memory can be read out from the memory and its mirror image simultaneously.

[0016] In some mirror-image embodiments, no memory allows pipelined read or write operations, for example, each 10 memory is an asynchronous SRAM. Other embodiments allow pipelined reads or writes or both.

[0017] The invention allows a non-blocking ATM switch performing multiple accesses to a shared memory in parallel to be easily constructed from switch fabrics each of which performs at most one shared-memory access at any given time.

[0018] The invention is not limited to ATM switches or to networks.

15 [0019] The invention is described further hereinafter, by way of example only, with reference to and as illustrated in the accompanying drawings in which:

Fig. 1 is a block diagram of an ATM switch having a memory system embodying the present invention.

Figs. 2A, 2B, and 2C are timing diagrams illustrating a switch cycle in some embodiments of Fig. 1.

20 Fig. 3 is a timing diagram of a synchronous SRAM used in some embodiments of Fig. 1.

Fig. 4 is a timing diagram illustrating read and write operations in some embodiments of Fig. 1.

Fig. 5 is a circuit diagram of a memory address generation logic in some embodiments of Fig. 1.

Fig. 6 is a block diagram showing some details of some embodiments of Fig. 1.

25 Fig. 7 is a block diagram of a portion of the ATM switch of some embodiments of Fig. 1.

Fig. 8 is a block diagram of a circuit used in some embodiments of Fig. 1.

Fig. 9 is a timing diagram of an address generation pipeline used in some embodiments of Fig. 1.

Fig. 10 is a block diagram of a circuit used in some embodiments of Fig. 1.

Fig. 11 is a block diagram of a portion of an ATM switch according to the present invention.

30 Fig. 12A, 12B are timing diagrams for the switch of Fig. 11.

Fig. 13 is a block diagram of an ATM switch according to the present invention.

Fig. 14 is a block diagram showing details of the switch of Fig. 13.

Fig. 15 is a block diagram of an ATM switch having a memory system according to the present invention.

35 Fig. 16 is a timing diagram of a switch cycle in some embodiments of Fig. 15.

[0020] Fig. 1 illustrates a high performance memory system including memory 110_UL ("UL" stands for "upper left", corresponding to the memory position in Fig. 1), memory 110_DL ("down left", or lower left), memory 110_UR ("upper right") and memory 110_DR ("down right", or lower right). These four memories 110 are single-ported. The four memories 110 share two address buses mAddr_U, mAddr_D and two data busses Data_L, Data_R. Address bus mAddr_U is connected to the address inputs of memories 110_UL and 110_UR. Address bus mAddr_D is connected to the 40 address inputs of memories 110_DL and 110_DR. Data bus Data_L is connected to the data ports of memories 110_UL and 110_DL. Data bus Data_R is connected to the data ports of memories 110_UR and 110_DR.

[0021] Every combination of an address bus and a data bus can be used to access one of the four memories. For example, address bus mAddr_U and data bus Data_L provide access to memory 110_UL; address bus mAddr_D and data bus Data_R provide access to memory 110_DR. Since there are two address buses and two data buses, two 45 memory accesses can proceed at the same time. In particular, a read and a write can proceed at the same time. Further, as explained below, because the memory system includes four memories, a read from any one of the four memories and a write to another one of the four memories can be performed every clock cycle even if the four memories use a different number of cycles for read and write operations.

[0022] Memories 110 form a shared memory system in ATM switch 118. Switch 118 is built from two switch fabrics 50 122.1, 122.2, and has twice the bandwidth and the number of ports of a single switch fabric. Each switch fabric 122 contains 64 ports (ports 0-63 in switch fabric 122.1, ports 64-127 in switch fabric 122.2). Switch 118 is non-blocking -- a cell received on any one of ports 0-127 can be transmitted to any one or more of ports 0-127.

[0023] Fig. 2A illustrates a switch cycle of switch 118. The switch cycle includes clock cycles 0-135. The clock cycle 55 numbers appear on top. Clock cycles 0-67 are the input phase of switch fabric 122.1 and the output phase of switch fabric 122.2. Clock cycles 68-135 are the output phase of switch fabric 122.1 and the input phase of switch fabric 122.2. In its input phase, the switch fabric reads ATM cell payloads from memories 110 to memory buffers 130 (the cell headers are stored in a different memory, not shown). In its output phase, the switch fabric writes cell payloads from buffers 130 to memories 110. Input and output phases of a switch fabric are described in U.S. Patent 5,440,523 issued August

8, 1995, and incorporated herein by reference, and PCT Publication No. WO 97/06489 published February 20, 1997, and incorporated herein by reference.

[0024] Operation during clock cycles 0-67 is illustrated in more detail in Fig. 2B. Cycles 68-135 are illustrated in more detail in Fig. 2C.

[0025] In clock cycle 0 (Fig. 2B), switch fabric 122.1 reads from one of memories 110 a cell payload to be transmitted on port 0, as indicated by "r0". In the same cycle, switch fabric 122.2 writes to another memory 110 a cell payload received on port 64, as indicated by "w64". Similarly, in clock cycle 1, a cell payload is read out for port 1 ("r1"), and a cell payload received on port 65 is written ("w65"). By the end of cycle 63, sixty-four cells (i.e. cell payloads) have been read out for ports 0-63, and the cells received on ports 64-127 have been written.

[0026] Switch fabric 122.1 includes four CPUs cpu0-cpu3 (not shown). Switch fabric 122.2 includes four CPUs cpu4-cpu7 (not shown). In cycles 64-67, switch fabric 122.1 reads from memories 110 four cells for transmission to respective CPUs cpu0-cpu3, and switch fabric 122.2 writes to memories 110 four cells received from respective CPUs cpu4-cpu7.

[0027] In cycles 68-135 (Fig. 2C), switch fabric 122.1 writes cells received from ports 0-63 and CPUs cpu0-cpu3 to memories 110. In the same cycles, switch fabric 122.2 reads cells to be transmitted to ports 64-127 and CPUs cpu4-cpu7.

[0028] A cell received by the switch can be written to any memory 110 for temporary storage before transmission. However, when the cell is to be transmitted, the cell has to be read from the memory in which it was stored. Therefore, for each clock cycle the switch 118 decides which of the four memories 110 has to be read during the clock cycle (the memory to be read is the memory storing the cell to be transmitted). Then the switch selects a different memory 110 to be written in the same clock cycle.

[0029] In some embodiments, memories 110 are synchronous SRAMs in which a read operation takes a different number of clock cycles from a write operation.

Synchronous SRAMs have the advantage of being fast. Further, the read operations are pipelined. However, differences between the read and write operation timing lead to a memory bus utilization penalty as illustrated in Fig. 3. In Fig. 3, a read operation takes two clock cycles and a write operation takes one clock cycle. In clock cycle 1 in Fig. 3, a write operation Wr1 is performed. The address is written to the address inputs A of the memory, and the data is written to the data port D. Then a read operation Rd2 is performed in cycles 2 and 3, and another read operation Rd3 is performed in cycles 3 and 4. More particularly, in cycle 2, the Rd2 address is supplied to the address inputs. In cycle 3, the Rd2 data are read out to the data port and the Rd3 address is supplied on the address inputs. In cycle 4, the Rd3 data are read out to the data port. In cycle 5, another write Wr4 is performed. The memory of Fig. 3 is part number MT 58LC64K36/B3 or MT 58LC128K18/B3 or MT 58LC64K18/B3 manufactured by Micron Corporation of Idaho. The memory is used in flow-through mode.

[0030] In Fig. 3, switching from a read operation to a write operation or from a write to a read involves a one-cycle penalty. With respect to the data port, the one-cycle penalty is seen between the write operation Wr1 and the read operation Rd2. With respect to the address inputs, the penalty is seen between the read operation Rd2 and the write operation Wr4.

[0031] The memory system of Fig. 1 avoids this penalty while using only two address buses and only two data buses. See the timing diagram of Fig. 4 and, in particular, the portion Fig. 4A of Fig. 4. In Fig. 4A, "Rdi" (i=1,2,3,4,5) indicates a read operation started in clock cycle i, and "Wri" indicates a write operation performed in cycle i. Successive reads from the same or different memories are pipelined. More particularly, while data from one read operation are provided on a data bus Data_L or Data_R, an address for another read operation is provided on address bus mAddr_U or mAddr_D. Hence, the read operations take one address bus and one data bus in every clock cycle. For example, in cycle 2, read operations Rd2 and Rd1 take address bus mAddr_U and data bus Data_L.

Therefore, the remaining address bus and the remaining data bus are available for a write. Thus, in cycle 2, address bus mAddr_D and data bus Data_R are used by write operation Wr2.

[0032] In some embodiments of Figs. 2B, 2C, "ri" (i = 0, 1, ... 127) or "rcpui" (i = 0, 1,...7) indicates a cycle in which a read address is provided on an address bus to read a cell to be transmitted respectively to port i or CPU cpui. In other embodiments, "ri" or "rcpui" indicates a cycle in which a cell to be transmitted to port i or CPU cpui is read out to a data bus.

[0033] The use of memories 110 in ATM switch 118 will now be described in more detail. As shown in Fig. 1, the switch includes port interface units (PIFs) 140.1 for ports 0-63 and PIFs 140.2 for ports 64-127. Interconnect matrix 150 connects PIFs 140 to memory buffers 130, as described in U.S. Patent 5,440,523 issued to A. Joffe on August 8, 1995 and incorporated herein by reference, and in PCT publication WO 97/06489 published February 20, 1997 and incorporated herein by reference. In switch fabric 122.1, switch fabric control logic 160.1 controls PIFs 140.1 and memory buffers 130. In identical switch fabric 122.2, switch fabric control logic 160.2 controls PIFs 140.2 and memory buffers 130. In addition, switch fabrics 122.1, 122.2 include cell queuing and scheduling logic which is not shown in Fig. 1. Such logic is described, for example, in the following U.S. Patent applications incorporated

herein by reference: serial number 08/706,104 filed August 30, 1996 by A. Joffe et al.; serial number 08/708,140 filed August 27, 1996 by A. Joffe; and serial number 08/845,710 filed April 25, 1997 by A. Joffe et al. See also the following publications incorporated herein by reference: "ATMS2000 User's Guide (MMC Networks, Inc., issue 1.1); "ATMS2004B Switch Controller 2 'GRAY'" (MMC Networks, Inc., document 95-0004); "ATMS2003B Switch Controller 1 'WHITE'" (MMC Networks, Inc., document 95-0003); "ATMS2000 Application Note: Support For More Than 32K Connections" (MMC Networks, Inc., document 95-0009).

[0034] Switch fabric control logic circuits 160.1, 160.2 are connected to control bus 164. Bus 164 is connected to address generation logic 170 (Figs. 1, 5). Logic 170 generates address signals for memories 110 and provides the address signals to buses mAddr_U, mAddr_D. Logic 170 also generates memory control signals described below.

[0035] In Fig. 5, OCP (Output Cell Pointer) is the read address, i.e., the address of the cell to be read. ICP (Input Cell Pointer) is the write address. Internally, the switch fabrics 122 supplement each memory 110 address by two least significant bits (LSBs) identifying a particular memory 110. Thus, in some embodiments, each memory 110 uses 15-bit addresses, and each address bus mAddr_U, mAddr_D is 15 bits wide, but the switch fabrics extend the addresses internally to 17 bits. In read addresses, the two LSBs identifying the memory 110 are referred to as the Output Cell Region (OCR).

[0036] OCP and ICP do not include the two LSBs identifying the memory 110. Thus, each of OCP and ICP is as wide as memory bus mAddr_U or mAddr_D (15 bits in the above example).

[0037] ICP, OCP and OCR are provided to control bus 164 by switch fabric control logic circuits 160 as described below. OCP is latched in latch 510. ICP is latched in latch 514. OCR is latched in latch 520. The output of latch 520 is connected to the input of logic circuit 524. Logic 524 generates select signals for multiplexers 530, 534. The two data inputs of each multiplexer are connected to respective outputs of latches 510, 514. In any given clock cycle, one of the multiplexers selects ICP from latch 514, and the other multiplexer selects OCP from latch 510. The output of multiplexer 530 is connected to bus mAddr_U through serially connected latches 540, 544. The output of multiplexer 534 is connected to bus mAddr_D through serially connected latches 550, 554.

[0038] Logic 524 operates in accordance with the algorithm described above in connection with Fig. 4A.

[0039] Logic 524 also generates the following signals:

- 1) write-enable signal L_we for memories 110_UL, 110_DL (Fig. 6);
- 2) write-enable signal R_we for memories 110_UR, 110_DR;
- 3) output-enable signal L_oe for memories 110_UL, 110_DL;
- 4) output-enable signal R_oe for memories 110_UR, 110_DR;
- 5) chip select signals UL_cs, UR_cs, DL_cs, DR_cs for respective memories 110_UL, 110_UR, 110_DL, 110_DR.

[0040] The timing diagram for these signals is illustrated in the Fig. 4B portion of Fig. 4. In Fig. 4B, "1" means the corresponding signal is asserted, and "0" means the signal is deasserted. (For example, in cycle 2, R_we is asserted and L_we is deasserted.) The embodiment being described uses memories in which for each memory access operation the chip select is asserted during the address phase of the operation, that is, during the clock cycle in which the address is provided to the memory. Similarly, the output enable signal is asserted during the address phase of a read operation, and the write enable signal is asserted during the write operation.

[0041] Logic 524 also generates two-bit signal Stk_Sel[1:0] which identifies the memory 110 that is to be written in a subsequent write operation, as described in more detail below.

[0042] Fig. 7 illustrates additional details of switch fabric control logic circuits 160.1, 160.2, control bus 164, and address generation logic 170.

[0043] Control bus 164 includes output control bus 164.0, input control bus 164.1, and a two-bit bus Stk_Sel connected to output Stk_Sel of logic 524 (Fig. 5).

[0044] Each fabric control logic 160.1, 160.2 includes four identical control blocks 710. Each control block 710 handles communication with eight of the ports 0-127 and with one of CPUs cpu0-cpu7. Thus, switch 118 is composed of eight switch fabrics, each of which is controlled by a respective control block 710. Each control block 710 has a modular switch controller (MSC) 720 connected to control bus 164 through control bus logic ("CBL") 730. Control bus logic 730 includes latches 740, 750 and queues 754, 758, 760. An output 770 of MSC 720 is connected through latches 740 to the following lines of output bus 164.0: MASK, SMID (Source Module ID), OCP, CID (Connection ID), and OCR. The MASK lines (the MASK "bus") include 1 bit for each MSC 720. The bit is set when the data on the output control bus 164.0 is to be read by the corresponding MSC 720.

[0045] Output 770 is also connected through latches 740 to the following lines of input bus 164.1: MASK, CID, SMID, ICP, and ATTR. ATTR is cell attributes. See "ATMS2004B Switch Controller 2 'GRAY'" cited above, pages 12-13.

[0046] Lines CID, SMID, ICP of input control bus 164.1 are connected to the input of queue 754 whose output is connected to MSC 720. Bus ATTR is connected to the input of queue 758 whose output is connected to MSC 720. Lines SMID, OCP, CID, OCR of output control bus 164.0 are connected to the input of queue 760 whose output is

connected to MSC 720. The MASK lines of buses 164.O, 164.I are connected to inputs of latch 750 whose output is connected to MSC 720.

[0047] Bus Stk_Sel is connected to each MSC 720 through respective CBL 730.

[0048] Lines ICP, OCP, OCR of control bus 164 are connected to inputs of address generation logic 170 as shown in Fig. 5.

[0049] As shown in Fig. 8, each MSC 720 stores four stacks 810_DL, 810_UL, 810_DR, 810_UR of 15-bit addresses (i.e. without the two LSBs identifying a memory 110) of all the free memory locations in respective memories 110_DL, 110_UL, 110_DR, 110_UR. Each memory location can store a single cell payload. In each clock cycle, the signal Stk_Sel causes multiplexers 820 of all the MSCs to select the stack corresponding to a memory 110 which will be written six cycles later. (As shown in Fig. 9, in clock cycle t-6 the Stk Sel signal selects the memory that will be written in cycle t.) The pointer popped from the selected stack is ICP. ICP is provided from the outputs of multiplexers 820 to the ICP lines of input bus 164.I through CBLs 730 in cycle t-3 (Fig. 9). In an earlier cycle, the MSC 720 corresponding to the port or CPU whose cell will be written in cycle t provides a MASK signal identifying all the MSCs 720 corresponding to the ports and CPUs to which the cell is to be transmitted. The MASK signal is written to the MASK lines of input bus 164.I through the respective latch 740 in cycle t-6.

[0050] When the cell has been transmitted to all the ports and CPUs, a pointer to the cell payload in a respective memory 110 is returned to a respective stack 810 (Fig. 8) in all MSCs 720.

[0051] Fig. 9 illustrates the address generation pipeline. In clock cycle t, one of the address buses mAddr_U, mAddr_D is driven with a read address. The other one of the address buses is driven with a write address. Some "n" cycles earlier (where n > 6), i.e. in cycle t-n, one of MSCs 720 provides, on the OCR bus, the OCR of the cell to be read in cycles t, t+1. In cycle t-6, address generation logic 170 generated the signal Stk_Sel identifying the memory 110 to be written in cycle t. Signal Stk_Sel is generated from the OCR signals driven on the OCR bus in cycles t-n and t-n-1. The OCR bus in cycle t-n identifies the memory that will receive a read address in cycle t. The OCR bus in cycle t-n-1 identifies the memory that will receive a read address in cycle t-1, and, therefore, identifies the data bus taken by the read operation in cycle t. See Fig. 4.

[0052] Also in cycle t-6, the control block 710 which is to write a cell in cycle t to memories 110 drives the MASK lines of output bus 164.O with signals identifying all the MSCs 720 which are to transmit the cell. The control block 710 which is to read a cell in cycle t drives the MASK lines of input bus 164.I with a signal indicating whether the control block needs the cell to continue to be stored in the memory after the cell is read out. This information is used to determine if the cell memory can be freed.

[0053] In cycle t-3, the MSC 720 which is to read a cell starting in cycle t provides the cell address OCP on the OCP bus. The MSC 720 which is to write a cell in cycle t provides the cell address ICP on the ICP bus. In cycle t-2, the read and write addresses appear respectively on the outputs of latches 510, 514 (Fig. 5). In cycle t, the addresses appear on buses mAddr_U, mAddr_D.

[0054] Returning to Fig. 1, the memory buffers 130, the interconnect matrix 150, and the PIFs 140.1 and 140.2 are similar to those described in the aforementioned U.S. patent 5,440,523 and PCT Publication WO 97/06489.

[0055] In some embodiments, interconnect matrix 150 includes two interconnect matrices, one for ports 0-63 and one for ports 64-127. Memory buffers 130 include two sets of memory buffers, one set for ports 0-63 and one set for ports 64-127. The interconnect matrix and the memory buffer set that correspond to ports 0-63 are part of switch fabric 122.1. The other interconnect matrix and the other memory buffer set are part of switch fabric 122.2. In each clock cycle, one set of the memory buffers is connected to data bus Data_L and the other set of the memory buffers is connected to data bus Data_R.

[0056] In other embodiments, a single interconnect matrix 150 and a single set of memory buffers 130 are shared by all the ports 0-127. Memory buffers 130 include 12 identical chips (i.e. integrated circuits) 130.1 through 130.12; see Fig. 10 showing a representative chip 130.i. A cell payload is transferred between memory buffers 130 and a PIF 140 in 12-bit words. See the aforementioned U.S. Patent 5,440,523 and publication WO 97/06489. Each chip 130.i stores the respective bits i of all 12-bit words of all cell payloads read to or written from memories 110. Each memory buffer chip 130.i includes: (1) a memory buffer 130.i.1 for bits i of words transferred to and from ports 0-63 and CPUs cpu0-cpu3, and (2) a memory buffer 130.i.2 for bits i of words transferred to or from ports 64-127 and CPUs cpu4-cpu7.

[0057] Each buffer 130.i.1, 130.i.2 is connected to interconnect matrix 150 by 68 one-bit lines. For each port and each CPU, the respective buffer 130.i.1 or 130.i.2 has: (1) a register (not shown) to hold bits i of 32 consecutive words received from that port or CPU, and (2) a separate register (not shown) to hold 32 bits i of 32 consecutive words to be transmitted to that port or CPU. In any given clock cycle, switch 1010 in chip 130.i connects one of buffers 130.i.1, 130.i.2 to data bus Data_R, and the other one of buffers 130.i.1, 130.i.2 to data bus Data_L. Each bus Data_R, Data_L is a 384-bit bus allowing a parallel transfer of a cell payload. (Of note, a cell payload includes exactly thirty-two 12-bit words, or 384 bits). All the switches 1010 of buffers 130.i are controlled by the same signal C generated by address generator 170 (Fig. 1). In any given clock cycle, all the buffers 130.i.1, for all i, are connected to one and the same of buses Data_L, Data_R, and all the buffers 130.i.2 are connected to the other one of the buses Data_L, Data_R. Thus,

one of switch fabrics 122.1, 122.2 uses one of the data buses, and the other switch fabric uses the other data bus.

[0058] Fig. 11 illustrates an alternate embodiment having separate address buses mAddr_UL, mAddr.UR, mAddr_DL, mAddr_DR for respective memories 110_UL, 110.UR, 110_DL, 110_DR. In each clock cycle, address generation logic 170 generates the same address on buses mAddr_UL, mAddr.UR, and also generates the same address on buses mAddr_DL, mAddr_DR. The remaining features of the system of Fig. 11 are similar to those of Fig. 1.

[0059] The separate address buses allow the switch of Fig. 11 to operate in an alternate mode in which different addresses can be generated on buses mAddr_UL, mAddr.UR, and on buses mAddr_DL, mAddr_DR. In this mode, each memory 110 has half the size and half the bandwidth of the Fig. 1 mode, and hence the memory system is less expensive. In particular, each data bus Data_L, Data_R has half the width of the respective data bus of the system of Fig. 1. Further, in the alternate mode, the switch has only 64 ports and four CPUs, which further reduces the system cost.

[0060] In the alternate mode, Fig. 2 still applies. Details of the input and output phases are illustrated in Figs. 12A, 12B. An example is given in Table 1 below. The ports are numbered 0, 2, 4, ..., 126 (odd numbers are omitted), and the four CPUs are labeled cpu0, cpu2, cpu4, cpu6. A read or a write of a single cell takes two clock cycles, with one half of a cell payload being read or written in each clock cycle. Thus, in clock cycle 0 (Fig. 12A), fabric 122.1 (Fig. 1) reads from one of memories 110 one half of a cell payload to be transmitted on port 0, as indicated by "r0". Fabric 122.1 reads the other half from the same memory in clock cycle 1, as indicated by "r0". Both halves of each cell payload are stored in the same memory to simplify keeping track of where the cell is. In cycles 0 and 1, switch fabric 122.2 writes to another memory 110 respective halves of a cell payload received on port 64 ("w64"). Similarly, in clock cycles 2 and 3, respective halves of a cell payload are read out for port 2 ("r2"), and respective halves of a cell payload received on port 66 are written ("w66"), and so on. In clock cycles 64, 65, fabric 122.1 reads a cell payload for transmission to cpu0, and fabric 122.2 writes a cell payload received from cpu4. In clock cycles 66, 67, fabric 122.1 reads a cell payload for transmission to cpu2, and fabric 122.2 writes a cell payload received from cpu6.

[0061] In clock cycles 68-135, fabrics 122.1, 122.2 exchange places, with fabric 122.1 being in the output phase and fabric 122.2 being in the input phase.

[0062] In Figs. 12A, 12B, "r" indicates a data phase of a read operation. Thus, in clock cycle 0, one half of a cell payload is provided on a data bus to fabric 122.1 in response to an address provided by logic 170 in the previous clock cycle. In clock cycle 1, the other half of the same cell payload is provided on the same data bus.

[0063] Table 1 below illustrates bus utilization in one example. In clock cycles 0, 1, read operation r0 reads memory 110_UL, taking the data bus Data_L. Data bus Data_R is available for the write operation w64. In cycle 0, the read r0 takes data bus mAddr_UL. In cycle 1, the next read operation r2 takes data bus mAddr_DR, to read data out of memory 110_DR in cycles 2 and 3. Therefore, address bus mAddr_DR is available for the write w64 in cycles 0 and 1.

Table 1

CLK	0	1	2	3	4	5
mAddr_UL	r0				w68	w68
mAddr_DR						
mAddr_DL			w66	w66		r6
mAddr_DR	w64	w64			r4	r4
Data_L	r0	r0	w66	w66	w68	w68
Data_R	w64	w64	r2	r2	r4	r4

[0064] We can see that in any two cycles $2i, 2i+1$ ($i = 0, 1, \dots, 67$), at least one of the memories 110 is available for a write operation. Indeed, without loss of generality, suppose the read operation in cycles $2i, 2i+1$ gets data from memory 110_UL. Since the read takes the data bus Data_L, the data bus Data_R is available for a write operation. The read takes bus mAddr_UL in cycle $2i$, and the next read will take one of the four address buses in cycle $2i+1$. Whichever address bus is taken, one of buses mAddr_DR, mAddr_DL remains available in cycle $2i+1$ for a write. Since the same bus is available for a write in cycle $2i$, at least one (and maybe both) of the memories 110_DR, 110_DL is available in both cycles $2i, 2i+1$ for two write operations, allowing both halves of a cell payload to be written into that memory in parallel with the reads.

[0065] As illustrated in Table 1, there is no data bus utilization penalty. However, two address buses are unused in every clock cycle in this mode.

[0066] In Fig. 13, the shared bus 164 of Fig. 7 is replaced by dedicated connections to reduce signal transmission noise at high frequencies such as 50MHz. Fig. 13 shows four cards (i.e., printed circuit boards) 1050.1, 1050.2, 1050.3, 1050.4. Each card 1050 includes two MSCs 720 and a control message interface ("CMI") integrated circuit 1054. The

MSCs and the CMI circuit are shown in card 1050.1 only for simplicity. One embodiment of CMI 1054 is the part CMI-50A available from MMC Networks of Sunnyvale, California.

[0067] Bus 164 is replaced by dedicated connections 164.i,j interconnecting respective cards 1050.i, 1050.j.

[0068] Fig. 14 illustrates a single representative card 1050 in more detail. Outputs 1060 of the two MSCs 720 are connected to a bus 1062 shared by the two MSCs. Bus 1062 is connected to the inputs of latches 1064.1, 1064.2, 1064.3 whose outputs are connected to backplane 164B. Backplane 164B includes the lines 164.i,j of Fig. 13. Each of latches 1064.1, 1064.2, 1064.3 provides an identical set of signals from bus 1062 to the backplane, which transmits the signals to respective three other cards 1050.

[0069] Bus 1062 is also connected to an input of CMI 1054 to allow the MSCs 720 to communicate with each other.

[0070] The backplane has lines connected to inputs of CMI 1054 through latches 1064. Outputs of CMI 1054 are connected to inputs of MSCs 720. These outputs perform the same function as the outputs of circuits 754, 758, 760, 750 of Fig. 7.

[0071] Abort logic 1070 has input/output ports connected to the two MSCs 720. Other input/output ports of abort logic 1070 are connected to backplane 164B through latches 1064 for transmission of the full signals OFULL, IFULL.

[0072] The signals going through backplane 164B are delayed by two levels of latches 1064 (one level on the source card 1050 and one level on the destination card). Therefore, CMI 1054 delays the signals received from bus 1062 of the same card (the loop-back signals) by two clock cycles.

[0073] In some embodiments, the circuitry of two cards 1050 (i.e., four MSCs 720 and two CMI circuits 1054) are combined on the same card. Other combinations are possible.

[0074] Fig. 15 illustrates another memory system in a non-blocking ATM switch 1110. Switch 1110 is composed of four switch fabrics 122.1-122.4. Each switch fabric 122 controls a number of ports (not shown), for example, 64 ports, and includes one or more CPUs (not shown), for example, four CPUs. The switch fabrics 122 are interconnected by a control bus (not shown). Switch 1110 also includes an address generation logic, memory buffers, PIFs, and one or more interconnect matrices, which are not shown for simplicity.

[0075] In Fig. 15, individual memories 110_A, 110_B, 110_C, 110_D, 110'_A, 110'_B, 110'_C, 110'_D are asynchronous SRAMs in some embodiments.

[0076] A switch cycle of switch 1110 is illustrated in Fig. 16. In clock cycles 0-67, switch fabrics 122.1, 122.2 are in the input phase, and fabrics 122.3, 122.4 are in the output phase. In clock cycles 68-135, fabrics 122.1, 122.2 are in the output phase, and fabrics 122.3 and 122.4 are in the input phase. In some embodiments, the input and output phases of each switch fabric are like those of Fig. 2.

[0077] The two switch fabrics that are in the input phase read memories 110 in every clock cycle. In any given clock cycle, the two input-phase fabrics may have to read data stored in the same memory. Therefore, in some embodiments, memories 110 are double-ported, or they are single ported but they are sufficiently fast to allow two read operations in a single clock cycle. However, in other embodiments, memories 110 are single-ported, and they are not sufficiently fast to allow two reads in a single clock cycle. Advantageously, these embodiments are less expensive. In these embodiments, each memory has a "mirror" image memory that contains the same data. Thus, memories 110_A, 110'_A are mirror images of each other, that is, they store the same data at any given time. Similarly, memories 110_B, 110'_B are mirror images of each other, memories 110_C, 110'_C are mirror images of each other, and memories 110_D and 110'_D are mirror images of each other. Hence, if two switch fabrics need data stored in the same memory in the same clock cycle, the two switch fabrics get the data from different memories that are mirror images of each other. For example, if switch fabrics 122.1, 122.2 both need data stored in memory 110_A, one of the two switch fabrics reads memory 110_A, and the other switch fabric reads memory 110'_A.

[0078] When any memory is written, its mirror-image memory is written with the same data at the same address. Thus, in one example, in one of clock cycles 68-135, switch fabric 122.3 writes a cell into each of memories 110_C, 110'_C at the same address for both memories, and switch fabric 122.4 writes another cell into each of memories 110_D, 110'_D at the same address for both memories. A double-ported memory is not needed for both memories.

[0079] Memory switch 1120 connects the data ports D of memories 110 to the data buses Data_1 through Data_4 of respective switch fabrics 122 as appropriate to accomplish the operation described above in connection with Fig. 16. Memory switch 1120 is controlled using methods known in the art. Address and control signals for memories 110 are generated using known methods.

[0080] Some embodiments combine the techniques of Figs. 1 and 13. For example, some embodiments include 32 single-ported synchronous SRAMs. The 32 SRAMs are subdivided into two blocks of 16 SRAMs each. The two blocks are mirror images of each other. We will denote the memories of one block as M_i,j where each of i, j takes the values 1, 2, 3 and 4. We will denote the memories of the other block as M'_i,j, where i, j take the same values. For each given i and j, the memories M_i,j and M'_i,j are mirror images of each other.

[0081] For each memory block, four data buses and four address buses are provided. For each given j, data bus Data_j is connected to the data ports of memories M_1,j, M_2,j, M_3,j, M_4,j. For each given i, address bus mAddr_i is connected to the address inputs of memories M_i,1, M_i,2, M_i,3, M_i,4. Similarly, for each given j, data bus Data'_j

is connected to the data ports of memories M'_i,j for all i , and address bus $mAddr'_i$ is connected to the address inputs of all the memories M'_i,j for all j .

[0082] The memory system is used in a non-blocking switch, for example, a non-blocking ATM switch, including four switch fabrics. In any given clock cycle, two of the switch fabrics are in the input phase and two of the switch fabrics are in the output phase, as in Fig. 16. If data from the same memory are desired by different switch fabrics at the same time, the switch fabrics read mirror image memories, as in Fig. 15. In each clock cycle, two address buses are taken by the read operations started in the current clock cycle and two data buses are taken by the read operations started in the previous clock cycle. See Fig. 4. Therefore, in each clock cycle, at least two pairs of "mirror" address buses ($mAddr_i1, mAddr'_i1$), ($mAddr_i2, mAddr'_i2$), and at least two pairs of "mirror" data buses ($Data_j1, Data'_j1$), ($Data_j2, Data'_j2$) are available for write operations. One of the two switch fabrics in the output phase writes a cell into "mirror image" memories $M_i1,j1, M'_i1,j1$, at the same address for both memories, and the other one of the switch fabrics in the output cycle writes a cell into memories $M_i2,j2, M'_i2,j2$, at the same address for both memories. Therefore, two reads and two writes are performed every clock cycle.

[0083] The above embodiments do not limit the invention. In particular, the invention is not limited to any particular kind of memories or memory timing. Some embodiments use DRAMs or other kinds of memories. In some embodiments of Fig. 1, each read operation takes more than two clock cycles and each write operation takes more than one clock cycle. In some embodiments, the number of address buses equals the number of clock cycles taken by a read operation, and a new read address is supplied on each clock cycle. The invention is not limited to the number of memories used in the memory system. In some embodiments, the memories that are mirror images of each other store the same data at the same address. In other embodiments, the different memories use different addresses to store the same data. In some embodiments, each memory is an integrated circuit or a combination of integrated circuits. The invention is not limited to ATM. Some embodiments include synchronous transfer mode or any other data flow. Some embodiments include Ethernet, frame relay, and/or PPP/Sonet switches. In some such embodiments, a packet or a frame is subdivided into fixed size (e.g., 48 byte) cells; the last cell may be padded to reach the fixed size. The cells are stored in memory as described above. Then the cells are read out of memory as described above, reassembled to obtain the original frame or packet, and the packet or frame is transmitted. Other embodiments and variations are within the scope of the invention, as defined by the appended claims.

30 Claims

1. A memory system comprising:

35 a plurality of memories (110);
 a plurality of address buses (Addr); and
 a plurality of data buses (Data),
 wherein at least one data bus (Data L) is connected to different memories (110UL, 110DL) of the said plurality of memories, which said different memories (110UL, 110DL) are connected to respective different address buses (AddrU, AddrD) of the said plurality of address buses.

40 2. The memory system of Claim 1, wherein each memory (110) is arranged to allow pipelined read operations in which the memory providing data in one read operation overlaps with the memory reading an address for another read operation.

45 3. The memory system of Claim 2, wherein each memory is a synchronous SRAM.

4. The memory system of Claim 1, and arranged to provide at least two overlapping read operations so that the provision of data on a data bus in one read operation overlaps with the provision of a read address on an address bus in another read operation, and

50 wherein the memory system allows the overlapping read operations to overlap with one or more write operations that use address and data buses not used by the overlapping read operations.

55 5. The memory system of any one of Claims 1-4, wherein for any combination of any address bus and any data bus there is a memory connected to the address bus and the data bus.

6. The memory system of any one of Claims 1-5, and arranged with a plurality of switch fabric circuits that are arranged to access the memories in parallel, wherein the number of the switch fabric circuits is equal to the number of the data buses.

7. A memory system comprising:

one or more memories;
 a plurality of address ports;
 5 a plurality of data ports such that a combination of an address port and a corresponding one of the data ports provides access to the memory system;
 a plurality of address buses connected to the address ports; and
 a plurality of data buses connected to the data ports,
 10 wherein at least one data bus is connected to different data ports corresponding to different address ports connected to different address buses.

8. The memory system of Claim 1, 5, 6 or 7, and arranged such that a read operation is allowed to overlap with a write operation such that a data bus utilization penalty is zero.

15 9. The memory system of Claim 7, wherein each combination of an address port and a corresponding data port are arranged for use for pipelined read operations in which provision of data from the data port in one read operation overlaps with the address port receiving an address for another read operation;

20 wherein successive write operations through any given address port and a corresponding data port are not pipelined; and
 wherein the memory system allows successive read operations to be performed in parallel with successive write operations, whether or not the successive read operations are performed through the same address port or the same data port and whether or not the successive write operations are performed through the same address port or the same data port.

25 10. The memory system of Claim 9, and arranged such that successive read operations and the successive write operations performed in parallel with the successive read operations can use all the address and data buses without any data bus utilization penalty.

30 11. The memory system of Claim 7, and arranged with a plurality of switch fabric circuits that are arranged to access the memory system in parallel wherein the number of the switch fabric circuits is equal to the number of the data buses.

35 12. A method of accessing a memory system comprising at least a first address port, a first data port, a second address port, and a second data port, the method comprising:

(1) in a read access through the first address and data ports, providing a read address signal to the first and second address ports;
 40 (2) in said read access, the first data port providing a read data signal; and
 (3) in parallel with at least a portion of step (2), in a write access through the second address port and the second data port, providing a write address signal to the first and second address ports and a write data signal to the second data port.

45 13. The method of Claim 13, wherein the first and second address ports are connected to a common line.

14. A method for accessing a memory system comprising at least a first address port, a first data port, a second address port, and a second data port, the method comprising:

50 (1) in a read access through the first address and data ports, providing a read address signal on a first line connected to the first address port;
 (2) in parallel with at least a portion of step (1), in a write access through the second address and data ports, providing a write address signal on a second line connected to the second address port and providing a write data signal on a third line connected to the first and second data ports; and
 55 (3) after providing the write data signal on the third line in step (2), the memory system providing a read data signal on the third line in said read access.

15. The method of Claim 12 or 14, wherein the first address and data ports are address and data ports of a first memory, and the second address and data ports are address and data ports of a second memory.

16. A memory system comprising:

a plurality of sets of memories; and
5 a circuit arranged for parallel accesses to the memories, such that:
when any memory in a set is written, all memories of the same set are written with the same data; and
a read of any memory in a set is arranged to be performed in parallel with a read of any other memory of the same set.

10 17. The memory system of Claim 16, further comprising a plurality of switch fabrics that are allowed access to the memory system.

18. The combination of Claim 17, wherein any switch fabric is allowed access to any memory of the memory system, so that the switch fabrics form a non-blocking switch.

15 19. A method of accessing a memory system comprising a plurality of sets of memories, wherein each set has at least M memories where $M > 1$, the method comprising:

(1) reading M of the memories in parallel; and
20 (2) in parallel with at least a portion of step (1), writing each memory of a memory set such that each memory of the memory set is written with the same data, wherein none of the memories being written is a memory read in step (1).

25

30

35

40

45

50

55




FIG. 1





FIG. 2A






FIG. 2C





FIG. 3

CLK	1	2	3	4	5	
mAddr_U	Rd1	Rd2	Wr3	Rd4	Wr5	
mAddr_D		Wr2	Rd3	Wr4	Rd5	
Data_L		Rd1	Wr3	Wr4	Rd4	Rd5
Data_R		Wr2	Rd2	Rd3	Wr5	
L_we		0	1	1	0	
R_we		1	0	0	1	
L_oe	1	0	0	1	1	
R_oe		1	1	0	0	
UL_cs	1	0	1	1	0	
UR_cs		1	0	0	1	
DL_cs	0	0	0	1	1	
DR_cs		1	1	0	0	

FIG. 4A

FIG. 4B

FIG. 4



170 → FIG. 5





FIG. 6



7




FIG. 8




FIG. 9




FIG. 10





FIG. 11







FIG. 13






FIG. 14




FIG. 15




FIG. 16

This Page Blank (uspto)



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 913 828 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
22.12.1999 Bulletin 1999/51

(51) Int Cl. 6: G11C 7/00

(43) Date of publication A2:
06.05.1999 Bulletin 1999/18

(21) Application number: 98308387.4

(22) Date of filing: 14.10.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 28.10.1997 US 959636

(71) Applicant: MMC Networks, Inc.
Sunnyvale, CA 94086 (US)

(72) Inventors:

- Joffe, Alexander
Palo Alto, CA 94306 (US)
- Birger, Ari
Sunnyvale, CA 94087 (US)

(74) Representative: Gill, David Alan
W.P. Thompson & Co.,
Celcon House,
289-293 High Holborn
London WC1V 7HU (GB)

(54) Memory system and method of accessing the same

(57) In a memory system, each data bus (data) is connected to memories (110) connected to different address buses (Addr). Each memory (110) allows pipelined read operations such that when data are being read out from a memory (110) in one read operation, the address can be provided to the memory (110) for another read. However, write operations are not pipelined, and the write address and write data are provided to the memory simultaneously. Nevertheless, consecutive reads can overlap with writes. Each write operation uses address (Addr) and data (Data) buses not taken by any read occurring in parallel with the write. The address

(Addr) and data (Data) buses are connected to the memories so that no data bus penalty occurs when a memory is switched from a read to a write or from a write to a read. In some embodiments, multiple memories are subdivided into sets of mirror-image memories. In each set, all the memories store the same data. When simultaneous read accesses are desired to read data stored in one of the memories, the read accesses can be performed instead to different memories that are mirror images of each other. When any memory is written, all the memories of the same set are written with the same data.




FIG. 1



**European Patent
Office**

EUROPEAN SEARCH REPORT

Application Number

EP 98 30 8387

DOCUMENTS CONSIDERED TO BE RELEVANT					
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)		
A	EP 0 201 174 A (RAYTHEON CO) 12 November 1986 (1986-11-12) * abstract; figures 2,3 *	1,7,12, 14,16,19	G11C7/00		
D,A	WO 97 06489 A (MMC NETWORKS INC) 20 February 1997 (1997-02-20) * figures 4,10,22 *	6,11,17, 18			
TECHNICAL FIELDS SEARCHED (Int.Cl.6)					
G11C G06F					
The present search report has been drawn up for all claims					
Place of search	Date of completion of the search	Examiner			
THE HAGUE	29 October 1999	Wolff, N			
CATEGORY OF CITED DOCUMENTS					
X : particularly relevant if taken alone	T : theory or principle underlying the invention				
Y : particularly relevant if combined with another document of the same category	E : earlier patent document, but published on, or after the filing date				
A : technological background	D : document cited in the application				
O : non-written disclosure	L : document cited for other reasons				
P : intermediate document	& : member of the same patent family, corresponding document				

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 98 30 8387

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

29-10-1999

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP 0201174 A	12-11-1986	AU	3478789 A	07-09-1989
		AU	582632 B	06-04-1989
		CA	1250370 A,C	21-02-1989
		CA	1264093 A	27-12-1989
		JP	1995109 C	22-11-1995
		JP	7031629 B	10-04-1995
		JP	61231639 A	15-10-1986
		US	4819152 A	04-04-1989
-----	-----	-----	-----	-----
WO 9706489 A	20-02-1997	US	5732041 A	24-03-1998
		AU	6590596 A	05-03-1997
		EP	0843854 A	27-05-1998
		JP	11510285 T	07-09-1999
		US	5910928 A	08-06-1999
-----	-----	-----	-----	-----

EPO FORM P059
For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

This Page Blank (uspto)